

# Evaluation of Computer Vision Models on Car Crash Detection

Navaneeth Dontuboyina

*Received September 20, 2023*

*Accepted February 07, 2024*

*Electronic access February 15, 2024*

This paper focuses on the Car Crash Dataset (CCD) which consists of dashcam footage capturing car crashes and non-crash scenarios. The dataset is derived from the Berkeley Deep Drive dataset (BDD100K) and includes annotations indicating crash presence in each frame. Specifically, the study emphasizes videos in which the accident is recorded from the driver's point of view, resulting in 801 videos with such involvement. First step was to organize and segment the videos, and subsequently filter them to retain footage filmed from the car's point view (or known in the dataset as ego-vehicles). Differing from previous studies that also tried to predict crashes between other, non-ego vehicles, this study is only focused on crashes between ego-vehicles and other objects because it is a logical first step for any autonomous vehicle to immediately stop if it senses it will crash. From there, research can be expanded to include crashes between other objects. This study does not use bounding boxes or other methods to isolate different components or objects within the driving footage. There are two approaches for modeling. The first employed pre-trained Long Short-Term Memory (LSTM) models such as InceptionV3, Visual Geometry Group 16 (VGG16), Residual Network 50 (ResNet50), chosen for their efficiency in image datasets and computer vision. The dataset was refined to exclude videos lacking ego-vehicle involvement. The data was then split for training and testing, and label processing was performed. Each model underwent feature extraction and utilized a Recurrent Neural Network (RNN) sequence model for sequential prediction. The second approach entailed a custom Convolutional Neural Network (CNN) model, differing from the LSTM framework. Results were documented in comma separated values (csv) files, and 40 non-crash observations were randomly selected to balance the pre-existing 40 crash observations. Threshold analysis was conducted using receiver operating characteristic (ROC) curves, determining optimal values for crash classification. The InceptionV3 model achieved the highest accuracy at 93%, with 2 false positives and 3 false negatives. VGG16 and ResNet50 achieved 87% and 88% accuracy, respectively, with 3 false positives and 7 false negatives. The custom model proposed in this research paper has exhibited 64% accuracy, with 19 false positives and 10 false negatives. The findings indicate promise in crash prediction for autonomous vehicles, with pretrained models showing robust performance. However, the custom model displayed lower accuracy and higher false negatives. There are, obviously, limitations to this paper that include a very general custom model, not much depth into the use of V2X technologies for car crash prediction and the use of only one camera angle as opposed to multiple angles from the car's perspective. All of these improvements are being discussed for further research in a follow-up paper. This study introduces a practical approach for autonomous vehicles to identify the primary threat of a front-end collision without relying on CCTV or similar footage, as seen in previous research. It also delves into how and if pre-trained LSTM models can be integrated into the complex architecture of computer vision models in the future.

## Introduction

The purpose of this research paper is to see whether or not simple, pre-trained LSTM keras models could sufficiently predict car crash detection and thus could be used as a part of complex CV model architecture. There have already been projects and studies on using computer vision to predict potential car crashes and aid with autonomous vehicles. Two projects/studies conducted by the same team of Dhananjai Chand, Savyasachi Gupta and Goutham K. at National Institute of Technology (NIT) Warangal, India give insight into this with two different types of ego-vehicle footage<sup>1,2</sup>. "Computer Vision-based Accident Detection in Traffic Surveillance" uses closed-circuit television (CCTV) footage from a variety of lighting and weather conditions in high urban intersections in India. The framework for the

machine learning algorithm utilizes a Mask R-CNN<sup>3</sup> which is used specifically for its region of interests (RoIs) align feature. This model also uses a box offset regressor, softmax classifier and a mask fully convolutional network (FCN) predictor. They classify the cars using bounding boxes and then try to predict around 15 frames ahead the likelihood of the bounding boxes hitting each other by assessing their vector angles and the acceleration of these objects. This model has a 71% detection rate and a 0.53% false alarm rate. "Computer Vision Based on Accident Detection for Autonomous Vehicles" is more related to this research topic as it captures the similar type of data which is crash footage from the dashcam and tries to achieve the similar sort of application for a self-driving car. It uses roughly a similar model as before with the same technique of drawing binding boxes and calculating the likelihood of a crash a few frames



```

path = "/Users/nalthan/Desktop/vertiasopencvprojecy/Crash-1500"

crash_file_names = [path + "/" + filename for filename in os.listdir(path) if filename.endswith('.mp4')]
print(crash_file_names)
for vid in crash_file_names:
    cam = cv2.VideoCapture(vid)
    currentframe = 1
    tempdir = "/Users/nalthan/Desktop/vertiasopencvprojecy/crashframes/"
    crashframedir = vid.replace("Crash-1500", "crashframes").replace(".mp4", '')
    os.mkdir(crashframedir)
    while(True):
        # reading from frame
        ret,frame = cam.read()

        if ret:
            # if video is still left continue creating images
            name = crashframedir + "/frame" + str(currentframe) + '.jpg'
            print ('Creating...' + name)

            # writing the extracted images
            cv2.imwrite(name, frame)

            # increasing counter so that it will
            # show how many frames are created
            currentframe += 1
        else:
            break
    cam.release()
    cv2.destroyAllWindows()

```

Fig. 2 Creating individuals frames for each video which has a car crash in it

```

path = "/Users/nalthan/Desktop/vertiasopencvprojecy/Normal"

normal_file_names = [path + "/" + filename for filename in os.listdir(path) if filename.endswith('.mp4')]
print(normal_file_names)
for vid in normal_file_names:
    cam = cv2.VideoCapture(vid)
    currentframe = 1
    tempdir = "/Users/nalthan/Desktop/vertiasopencvprojecy/normalframes/"
    normalframedir = vid.replace("Normal", "normalframes").replace(".mp4", '')
    os.mkdir(normalframedir)
    while(True):
        # reading from frame
        ret,frame = cam.read()

        if ret:
            # if video is still left continue creating images
            name = normalframedir + "/frame" + str(currentframe) + '.jpg'
            print ('Creating...' + name)

            # writing the extracted images
            cv2.imwrite(name, frame)

            # increasing counter so that it will
            # show how many frames are created
            currentframe += 1
        else:
            break
    cam.release()
    cv2.destroyAllWindows()

```

Fig. 3 Creating individuals frames for each video which does not have a car crash in it

```
def build_feature_extractor():
    feature_extractor = tf.keras.applications.ResNet50(
        include_top=True,
        weights="imagenet",
        input_tensor=None,
        input_shape=None,
        pooling=None,
        classes=1000,
    )

    preprocess_input = tf.keras.applications.resnet.preprocess_input

    inputs = keras.Input((IMG_SIZE, IMG_SIZE, 3))
    preprocessed = preprocess_input(inputs)

    outputs = feature_extractor(preprocessed)
    return keras.Model(inputs, outputs, name="feature_extractor")

feature_extractor = build_feature_extractor()
```

Fig. 4 Sample feature extractor for ResNet50 model

```
# For each video.
for idx, path in enumerate(video_paths):
    print(idx, path)
    print(type(path))
    # Gather all its frames and add a batch dimension.
    frames = load_video(os.path.join(root_dir, str(path).zfill(6)+'.mp4'))
    frames = frames[None, ...]

    # Initialize placeholders to store the masks and features of the current video.
    temp_frame_mask = np.zeros(shape=(1, MAX_SEQ_LENGTH, 1), dtype="bool")
    temp_frame_features = np.zeros(
        shape=(1, MAX_SEQ_LENGTH, NUM_FEATURES), dtype="float32"
    )

    # Extract features from the frames of the current video.
    for i, batch in enumerate(frames):
        video_length = batch.shape[0]
        length = min(MAX_SEQ_LENGTH, video_length)
        for j in range(length):
            temp_frame_features[i, j, :] = feature_extractor.predict(
                batch[None, j, :]
            )
            temp_frame_mask[i, :length] = 1 # 1 = not masked, 0 = masked

    frame_features[idx, :] = temp_frame_features.squeeze()
    frame_masks[idx, :] = temp_frame_mask.squeeze()

return (frame_features, frame_masks), labels

train_data, train_labels = prepare_all_videos(train_df, "/Users/nalthan/Desktop/vertiasopencvprojecy/train")
test_data, test_labels = prepare_all_videos(test_df, "/Users/nalthan/Desktop/vertiasopencvprojecy/test")
```

Fig. 5 ResNet50 feature extractor being applied to the the train and test data

a crash or non-crash using the ROC curves from sklearn kit python library. After determining the threshold value that yields the most accuracy on the test data for each of the models, a confusion matrix was made to see how well the model predicted for the test data.

## Results

The InceptionV3 model achieved 92% accuracy while reporting two false positives and three false negatives of whether a car crash occurs. The VGG16 model and ResNet50 model achieved 87% and 88% accuracy respectively while reporting three false positives and seven false negatives of whether a car crash occurs.

The custom model developed for this research achieved 64% accuracy while reporting 19 false positives and 10 false negatives of whether a car crash occurs.

	Thresholds	Accuracy
<b>8</b>	0.2958	0.938272
<b>10</b>	0.2827	0.925926
<b>9</b>	0.2879	0.913580
<b>12</b>	0.2343	0.876543
<b>13</b>	0.2246	0.876543

Table 1 The top 5 threshold values and their corresponding accuracy rates for the InceptionV3 model





Fig. 6 Confusion Matrix for Inception V3 model

	Thresholds	Accuracy
<b>8</b>	0.2370	0.876543
<b>9</b>	0.2289	0.876543
<b>6</b>	0.2471	0.876543
<b>7</b>	0.2421	0.864198
<b>4</b>	0.3011	0.839506

Table 2 The top 5 threshold values and their corresponding accuracy rates for the VGG16 model

## Discussion


These models have comparable results to previous, aforementioned studies and projects in predicting and identifying car



**Fig. 7** Confusion Matrix for VGG16 model

	Thresholds	Accuracy
<b>8</b>	0.1874	0.888889
<b>6</b>	0.1952	0.876543
<b>7</b>	0.1917	0.864198
<b>10</b>	0.1825	0.839506
<b>9</b>	0.1831	0.827160


**Table 3** The top 5 threshold values and their corresponding accuracy rates for the ResNet50 model



**Fig. 8** Confusion Matrix for ResNet50 model

	Thresholds	Accuracy
<b>23</b>	0.2163	0.641975
<b>12</b>	0.2433	0.641975
<b>16</b>	0.2292	0.629630
<b>25</b>	0.2149	0.629630
<b>24</b>	0.2157	0.629630

**Table 4** The top 5 threshold values and their corresponding accuracy rates for the custom model



**Fig. 9** Confusion Matrix for the custom model

crashes for autonomous vehicles. This study's accuracy rates of 93.8%, 87.6%, 88.8% and 64.1% for the LSTM and CNN custom model respectively are similar, if not better, to the NIT Warangal prediction rate of 71% and the VSLab rate of 79%. These results are probably because the past studies tried to use more advanced and complex methods such as bounding boxes and individual object detection as well as trying to account for more factors like crashes that do not include the self-car. The more experimental nature of these past studies result in lower prediction rate even though the methods are quite fascinating and innovative. However, it is important to note that the custom built model performed worse in terms of accuracy and minimizing false negatives. The reason why the InceptionV3 model did the best is most likely because it possesses the fastest GPU and CPU inference time of 42.2 ms and 6.9 ms respectively. This study presents an affirmative step into how an autonomous vehicle would realistically identify the most basic threat of a crash from the front rather than utilize CCTV or other types of footage as used in previous research. It also suggests that the process of identifying these basic threats is best when using a pre-trained computer vision model and combining that with

---

an RNN structure. These results mean that in the future when industries are developing CV models for car crash prediction and detection, they are predicted to use a simple LSTM that is most like the architecture of the InceptionV3 model due to its fast inference time and its formidable accuracy rate. Beyond autonomous vehicles, there are far more implications in other parts of the automotive industry based on this study. Potentially, traffic signals and city-wide traffic management systems can use these CV models to help in flow of traffic, or identify traffic violators. Insurance companies can implement CV models to identify who is at fault in an accident.

## Conclusions

The research has led to a conclusion that the InceptionV3 model achieved the highest accuracy at 93%, with 2 false positives and 3 false negatives. VGG16 and ResNet50 achieved 87% and 88% accuracy, respectively, with 3 false positives and 7 false negatives. The custom model proposed in this research paper has exhibited 64% accuracy, with 19 false positives and 10 false negatives. The findings indicate promise in crash prediction for autonomous vehicles, with pretrained models showing robust performance. However, the custom model displayed lower accuracy and higher false negatives. For further research, the next step is to most certainly incorporate real time decision making. A model that can take in a video and predict a certain amount of frames or seconds ahead as to whether or not a crash will occur lends itself more useful to self-driving cars. Then another addition would be to do more research into specifically building a custom model for predicting car crashes. The custom model used in this research was basic and more general which is why the accuracy paled in comparison to the pretrained models. Increasing the model to handle longer video lengths and inputs for videos from different angles of the car such as the back, side and blindspots for the same occurrence will be an improvement to look out for as this is how car manufacturers collect data for their autonomous vehicles. Another next step I would like to pursue to research more to improve predictability is on autonomous cars with Connected Vehicle (CV) technologies such as vehicle-to-everything (V2X) - vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I) and vehicle-to-pedestrian (V2P)<sup>8</sup>. These connected vehicle technologies would provide communication among vehicles in their nearby area and their surrounding environment. For a crash detection and prediction computer vision model to be more accurate, it must also take into account CV technology and their different V2X facets as different inputs.

## Acknowledgements

Thank you for the guidance of Mr. Nowell Closser from Harvard University in the development of this research paper.

## References

- 1 P. Ijjina, D. Chand, S. Gupta and K. Goutham, 2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT, Kanpur, India, pp. 1–6,.
- 2 S. Chand and I. Kavati, 2020 IEEE 17th India Council International Conference (INDICON, New Delhi, India, pp. 1–6,.
- 3 Kaiming, *Mask R-CNN*, arXiv.Org, 24 Jan. 2018, arxiv.org/abs/1703.06870.
- 4 Wentao, Proceedings of the 28th ACM International Conference on Multimedia.
- 5 of *Self-Driving Vehicles.” Coalition For Future Mobility, coalitionforfuturemobility.com/benefits-of-self-driving-vehicles/*, Accessed 1 Feb.
- 6 A. Jamil, *Kaggle*, ---.
- 7 Wentao, *GitHub*, 2012.
- 8 *U.S. Department of Transportation*, -----.